

Reprint Series**NURC-PR-2006-007**

Benchmark Evaluation of Multistatic Trackers

Coraluppi Stefano, Grimmett Doug, Pascal de Theije

August 2006

Originally published in:

FUSION 2006, The 9th International Conference on Information
Fusion 2006, Florence (Italy), 10-13 July 2006, Conference
Proceedings.

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE AUG 2006	2. REPORT TYPE	3. DATES COVERED 00-08-2006 to 00-08-2006			
4. TITLE AND SUBTITLE Benchmark Evaluation of Multistatic Trackers			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NATO Undersea Research Centre, Viale San Bartolomeo 400, 19138 La Spezia, Italy, ,			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

NATO Undersea Research Centre (NURC)

NURC conducts world class maritime research in support of NATO's operational and transformational requirements. Reporting to the Supreme Allied Commander Transformation, the Centre maintains extensive partnering to expand its research output, promote maritime innovation and foster more rapid implementation of research products.

The Scientific Programme of Work (SPOW) is the core of the Centre's activities and is organized into four Research Thrust Areas:

- Expeditionary Mine Countermeasures (MCM) and Port Protection (EMP)
- Reconnaissance, Surveillance and Undersea Networks (RSN)
- Expeditionary Operations Support (EOS)
- Command and Operational Support (COS)

NURC also provides services to other sponsors through the Supplementary Work Program (SWP). These activities are undertaken to accelerate implementation of new military capabilities for NATO and the Nations, to provide assistance to the Nations, and to ensure that the Centre's maritime capabilities are sustained in a fully productive and economic manner. Examples of supplementary work include ship chartering, military experimentation, collaborative work with or services to Nations and industry.

NURC's plans and operations are extensively and regularly reviewed by outside bodies including peer review of the research, independent national expert oversight, review of proposed deliverables by military user authorities, and independent business process certification. The Scientific Committee of National Representatives, membership of which is open to all NATO nations, provides scientific guidance to the Centre and the Supreme Allied Commander Transformation.



Copyright © NATO Undersea Research Centre 2005. NATO member nations have unlimited rights to use, modify, reproduce, release, perform, display or disclose these materials, and to authorize others to do so for government purposes. Any reproductions marked with this legend must also reproduce these markings. All other rights and uses except those permitted by copyright law are reserved by the copyright owner.

NOTE: The NURC Reprint series reprints papers and articles published by NURC authors in the open literature as an effort to widely disseminate NURC products. Users are encouraged to cite the original article where possible.

Benchmark Evaluation of Multistatic Trackers

Stefano Coraluppi, Doug Grimmett¹

NATO Undersea Research Centre (NURC)

Viale S. Bartolomeo 400, 19138 La Spezia, Italy

coraluppi@nurc.nato.int

Pascal de Theije

TNO Defence, Security, and Safety

P.O. Box 96864, 2509 JG The Hague

The Netherlands

pascal.detheije@tno.nl

Abstract - This paper provides an overview of the Special Session on Multistatic Sonar and Radar Tracking at FUSION 2006. This includes background on the Multistatic Tracking Working Group, a brief description of the datasets and trackers that compose this working group at present, and a detailed discussion of a proposed set of tracker performance metrics. We identify a number of issues associated with performance assessment for target tracking. We conclude with recommendations for continued performance assessment of multistatic trackers.

Keywords: Active Sonar - Multistatic Sonar and Radar - Sensor Fusion – Target Tracking – Performance Evaluation and Benchmarking.

1 Introduction

Anti-submarine warfare (ASW) operations are challenged due to the quiet nature of current threat submarines, and the complexity of shallow-water acoustic environments. Multistatic operations with low-frequency active sonar (LFAS)-equipped assets or deployable systems have the potential to improve ASW operations by exploiting detection information from a number of source-receiver combinations. The high data rate associated with a multistatic operation places added importance on data fusion and target tracking technology.

Multistatic sonar surveillance scenarios are generally based on one of two system concepts: mobile LFAS-equipped platforms (suitable for expeditionary tasks), and fixed/drifting deployed fields (suitable for surveillance of ports, harbors, choke points, etc.).

Both system concepts are under evaluation in several of the laboratories participating in the Multistatic Tracking Working Group. An example of the first system concept is illustrated in Figure 1, which shows the monostatic and bistatic source-receiver combinations that were used in a 2003 sonar sea trial jointly conducted by NURC and TNO. Figure 2 illustrates the second system concept: deployable sonar equipment used in a moored configuration.




Fig. 1. An example of a mobile surveillance network (bistatic detection in blue, monostatic in red).




Fig. 2. An example of a fixed surveillance network.

1.1 Fusion and Tracking

Numerous approaches to multi-sensor fusion and tracking are documented in the literature. Most of these follow one of two basic paradigms: contact-based, Kalman filter based approaches [1-3], and unified detection and tracking approaches [4]. Each of the well-known references that we indicate exhibits the particular bias of the authors for a distinct approach to the tracking problem:

- *Probabilistic data association* (PDA): scan-based, with soft data association [1];

¹ Work performed while at NURC. Current e-mail address: grimmett@spawar.navy.mil.

- *Probabilistic multi-hypothesis tracking* (PMHT): batch-processing based, with soft data association [2];
- *Multi-hypothesis tracking* (MHT): multi-scan based, with hard data association [3];
- *Bayesian tracking*: likelihood-surface based tracking, with either matched-filtered or contact-level inputs [4].

Each of these fundamental approaches has spawned its own literature, in which various enhancements are brought to bear on the tracking problem including the Interacting Multiple Model (IMM) filter, particle filters, the use of amplitude and classification information, etc. The literature provides a vast assortment of trackers; at best, authors generally provide performance assessment with respect to a simple tracker that utilizes another approach and for a selected set of metrics. Thus, when one asks a question like “Is the MHT tracker better than the PMHT?” the answer is another set of questions: “Which MHT? Which PMHT? For what data? With respect to what metrics?”

Clearly, an exhaustive evaluation of all variations of trackers based on the four paradigms listed above, for all relevant scenarios, and with respect to all possible metrics of interest, is impossible. Furthermore, it is difficult for one researcher or laboratory to be sufficiently knowledgeable and to have sufficient research focus to do justice to this task. Even a partial evaluation along these lines requires a team effort with participation from a number of laboratories.

In the radar tracking community benchmark problems have been defined, to which distinct tracking approaches have been applied [5]. To our knowledge, a similar effort does not exist in active sonar tracking.

1.2 Multistatic Tracking Working Group

In late 2004, the Multistatic Tracking Working Group (MSTWG) was set up. Defining characteristic of the MSTWG include the following:

- The intent of the working group is to foster the exchange of scientific and technical ideas, problems, and solutions related to multistatic tracking for sonar and radar.
- This will include the collaborative analysis of common data sets and will culminate in a workshop or special session disseminating final results towards the end of 2007.
- The group shall be comprised of delegates from various NATO nations and shall operate under a defined charter for a period of three years.
- Membership in the working group is limited to those within NATO nations who are actively researching or have expertise in multistatic

tracking. The working group has the authority to add members satisfying these restrictions.

The first and second meetings of the MSTWG took place at The Hague (April 2005) and Bonn (September 2005). The third meeting will be held in La Spezia (July 2006), and includes participation in this Special Session on Multistatic Sonar and Radar Tracking under the auspices of FUSION 2006 in Florence. The purpose of the session is to provide an overview of the MSTWG, to introduce the current common datasets for analysis, and to provide preliminary tracking performance results. More definitive performance results will be reported in 2007.

At present, membership of the MSTWG includes the following nations, laboratories, and delegates:

- NATO: Stefano Coraluppi (NURC), Doug Grimmett (NURC)²;
- NL: Mathieu Colin (TNO), Pascal de Theije (TNO), Leon Kester (TNO);
- GE: Frank Ehlers (FWG)³, Wolfgang Koch (FGAN);
- US: Brian La Cour (ARL/UT), Warren Fox (APL/UW), Christian Hempel (NUWC), James Pitton (APL/UW), Roy Streit (Metron), Peter Willett (UCONN).

The current membership of the MSTWG encourages the interest and potential participation in the work of the group by additional laboratories and delegates.

1.3 Outline of Paper

This paper is organized as follows. In section 2, we identify the datasets and trackers used by the MSTWG. In section 3, we define the input and output performance metrics of interest. Some of these metrics are not directly useable to evaluate some trackers; for example, those algorithms that operate with matched-filtered data do not have input detection statistics. These and other issues in tracker performance evaluation are discussed in section 4. Section 5 provides recommendations for future work by the MSTWG.

2 Datasets and Trackers

Numerous approaches exist for data simulation. All delegates in the MSTWG have agreed for their algorithms to track with detection-level (contact-level) data. Typically, this means that each (ping, source, receiver) triple in the multistatic network will generate on the order of hundreds of contacts for further processing. Contact-level data is particularly useful in multistatic (or more generally, multisensor) network-centric operations where

² Now with SPAWAR (USA).

³ Now with NURC.

real-time communication bandwidth limitations between platforms are of concern.

A typical processing chain with actual sea trial hydrophone data includes a number of steps as illustrated in Figure 3.




Fig. 3. An example of an active sonar processing chain.

2.1 MSTWG Datasets

Three simulation-based benchmark datasets have been generated thus far by the MSTWG for tracker analysis, based on the following distinct approaches:

1. Hydrophone level, time series simulation [6];
2. Contact-level, detection data simulation [7];
3. Hybrid simulation approach: real hydrophone-level environmental data, with injected, synthetic target data [8].

The first two scenarios are based on the mobile-platform system concept, as in Figure 1. The third leverages NURC environmental data collected in 2004 with a network of buoys as illustrated in Figure 2.

Thus far, only FM-based transmissions have been considered. By their nature, these provide only positional information on targets. The first and third simulation approaches require the use of processing chains similar to Figure 3.

In all cases, the simulations eventually result in a set of contact files, each of which contains receiver locations, array heading information, and a set of (time, bearing) contacts. Note that in real operations the generation of contact files requires that transmitted waveform information be available at the receiver processing chain, as well as knowledge of source location and ping time (with knowledge of sound speed, each of these can be estimated from the other).

This requires the exchange of this ancillary information, or meta-data (waveform, source location, and ping time), generally through radio links.

None of the benchmark datasets introduces systematic bias errors in any of the measured quantities; this will be addressed in future simulations. However, two of the three simulated datasets do include random errors in the fundamental measurements: sensor positions, sound speed, sensor heading, and time synchronization. Registration errors are a particular challenge for multisensor fusion [9]. In the literature, algorithms for automated estimation and removal of bias errors generally assume known data association; alternatively, algorithms for data association generally assume no bias errors (or small, residual errors that are treated as random measurement errors). The simultaneous solution to the bias-estimation and data-association problems is beyond the scope of the benchmark study undertaken to date. We assume that a sufficient amount of system calibration has been achieved, possibly coupled with a multistatic-specific registration solution like direct-blast stabilization.

2.2 MSTWG Trackers

The following tracking approaches are represented as part of our performance study:

- Parametric: ML-PDA & ML-PMHT [11];
- Soft data association: IMMPDAFAI [11], PMHT [12];
- Hard data association: GNN [13], MHT [14, 15], D-MHT [15];
- Bayesian [16].

Note that *parametric tracking* refers to the determination of target trajectory parameters; this approach has implicit single-target and fixed-order target-motion assumptions. Also, GNN (which denotes *global nearest neighbor* tracking) is the single-scan version of MHT [3]. Finally, D-MHT denotes *distributed MHT*, with scan-based track fusion [10].

3 Measures of Performance

In detection theory, it is common to use a *receiver operating characteristics* (ROC) curve as a complete statistical characterization of performance. For fusion and tracking systems, the situation is more complex: as noted in [3], “*There is no universally accepted set of tracking system figures of merit or MOEs.*”

For the purposes of this study, we have proposed a small set of metrics that we believe captures the salient features of tracker effectiveness for multistatic operations. The advantage of utilizing a common set of metrics is the opportunity for cross-evaluation of trackers. In general, we expect that each tracking approach will exhibit distinct strengths and weaknesses in a scenario-dependent manner.

Our input performance metrics assume contact-level input data. All contacts must be classified as *true* (i.e. target-originated) or *false*. Depending on the nature of the data

(i.e. real vs. simulated, and simulation methodology), these tags may require a pre-specified distance threshold to ascertain whether a contact is close enough to the target ground truth location for it to be deemed a target-originated detection. With real or hydrophone-level simulated data, the evaluation of input contact data quality is a function of this distance threshold.

A similar issue relates to the classification of output data: each track must be classified as *true* or *false*. One approach that is applicable to trackers that employ hard data association is, for each output track, to evaluate the fraction of target-originated contacts in the track. If this fraction is above a truth-acceptance threshold (say 0.5), then the track is classified as *true*; otherwise, it is classified as *false*. Alternatively, a distance metric can be used for track classification; this is applicable whether or not hard data association is used.

Once input contact data and output tracks have been truth-tagged, the input and output metrics defined below can be evaluated.

3.1 Tracker Input Metrics

- *Probability of detection (P_D)*: $P_D = \frac{N_C}{N_{CF} N_T}$, where N_C is the number of target-originated contacts, N_{CF} is the number of contact files, and N_T is the true number of targets.
- *False-alarm rate (FAR)*: $FAR = \frac{N_{NC}}{T}$ [s⁻¹], where N_{NC} is the number of non-target originated contacts, and T is the time duration of the scenario.
- *Localization error (LE)*: $LE = ave\left(\left\|\begin{bmatrix} x_C \\ y_C \end{bmatrix} - \begin{bmatrix} x \\ y \end{bmatrix}\right\|\right)$ [m], where $\begin{bmatrix} x_C \\ y_C \end{bmatrix}$ is a contact location, $\begin{bmatrix} x \\ y \end{bmatrix}$ is the corresponding true target location, $\|\cdot\|$ denotes the Euclidean norm, and the average is computed over all target-originated contacts.

3.2 Tracker Output Metrics

- *Track Probability of Detection (TP_D)*: $TP_D = \frac{\sum_{i=1}^{N_T} T_i}{T}$, where N_T is the number of true tracks, and T_i is the time duration of the i^{th} true track. TP_D is the ratio of the total duration of all true tracks and the total scenario duration. For each true track, the time duration is defined as the difference in time of the last and first contact that the track associates. Time overlaps in true tracks for the same target are removed.

- *Track false-alarm rate (TFAR)*: $TFAR = \frac{N_{FT}}{T}$ [s⁻¹], where N_{FT} is the number of false tracks.
- *Track-localization error (TLE)*: $TLE = ave\left(\left\|\begin{bmatrix} x_T \\ y_T \end{bmatrix} - \begin{bmatrix} x \\ y \end{bmatrix}\right\|\right)$ [m], where $\begin{bmatrix} x_T \\ y_T \end{bmatrix}$ is a track location, $\begin{bmatrix} x \\ y \end{bmatrix}$ is the corresponding true target location, $\|\cdot\|$ denotes the Euclidean norm, and the average is computed over all N_T true tracks at all ping times.
- *Track fragmentation rate (TFR)*: $TFR = \frac{N_T}{T \cdot N_T}$ [s⁻¹]. It may occur that the tracker algorithm fragments tracks. This means the tracker is unable to continuously output a single true track for the entire target trajectory. The tracker may lose the target, and subsequently reacquire it. This metric quantifies the average number of true tracks per target per unit time.
- *Latency (L)*: The tracker latency is the worst-case time lag [s] from input to output. For example, in multi-hypothesis tracking, the output may lag the input by a few scans of data. Other algorithms will have no lag. A batch algorithm that uses the entire dataset will have a lag equal to the scenario duration.
- *Execution rate (ER)*: Ratio of tracker execution time and scenario time; must be less than unity to achieve a real-time processing requirement.

3.3 Comments on Proposed Metrics

The choice of proposed metrics is motivated by the need to have a concise characterization of input data quality and of tracker performance. Furthermore, the first three tracker metrics are directly comparable to the input metrics, allowing for an assessment of *fusion gain*. Indeed, if one considers each input contact as a distinct track, one could assess the *fragmentation gain* as well, using the fourth tracker metric. Latency and execution rate are not directly linked to tracker input metrics and identify the timeliness of track information (the input latency may be viewed as being zero).

Generally, input detection performance is quantified over a range of detection thresholds, leading to a ROC performance curve (P_D vs. *FAR*). Similarly, output detection performance can be assessed over a range of detection thresholds and/or tracking parameters, with an output ROC curve (TP_D vs. *TFAR*).

It is possible to disregard latency as a metric and, instead, to reflect its impact in localization error computations. That is, by predicting track estimates to the current time

based on position and velocity information, one can assess current-time localization errors.

4 Issues in Performance Evaluation

There are several issues associated with tracker-performance evaluation. A number of these are identified below.

4.1 Track-to-Truth Assignment

One difficulty with the classification of tracks as *true* or *false* relates to short-duration tracks. These tracks are easily misclassified; for instance, even if a sequence of target-originated contacts leads to a confirmed track, only the confirmed portion of the track (that which is available at the output) is reported. Thus, a short confirmed track might easily be classified as *false*, if subsequent to confirmation the track includes a few false contacts. Likewise, a spurious short-duration track might easily be classified as *true*, if subsequent to confirmation the track includes a few target-originated contacts.

A second, more problematic difficulty associated with the classification of tracks relates to long-duration tracks. The issue is illustrated in Figure 4.




Fig. 4. Crossing targets with track swapping.

In Figure 4, we have an example of crossing targets, for which the tracker erroneously generates tracks that approach one another, and then diverge without crossing. To which ground truth is each track compared for *TLE* evaluation? It is possible to assign tracks to truth on a scan-by-scan basis. That is, a track may be mapped to one target for a portion of the run, and subsequently to another. This approach differs from the global track-to-truth assignment that we have chosen.

There are deficiencies in the scan-by-scan mapping approach. Indeed, it is important to penalize in some manner the track-swapping phenomenon, which can be extremely damaging operationally. With track-to-truth reassignment, this penalization is not directly achieved.

Similarly, a track that is close to a target for a portion of the run, and is associated with false contacts for another portion, leads to a difficult track-to-truth assignment problem. This is illustrated in Figure 5. This example illustrates a target track that is lured away by a region of fixed clutter returns. Correspondingly, the fixed clutter track is lured away by the target returns. In this instance, as above, it is difficult to classify each track as *true* or *false*.

For each track for which truth determination is difficult, one will either optimistically classify the track as *true*, or pessimistically classify it as *false*. In the former case, one will tend to overestimate TP_D , and to overestimate *TLE* as well. In the latter case, one will tend to underestimate TP_D and to underestimate *TLE* as well.




Fig. 5. True and false track confusion.

As noted previously, a simplified version of the *true* vs. *false* classification issue exists with input contact-level data. As the distance threshold for the classification of a contact as target-originated increases, the resulting ROC-curve performance increases, with a corresponding increase in *LE*. Thus, there is a tradeoff in our assessment of detection and localization statistics.




Fig. 6. Track fragmentation may lead to multiple tracks on the same target for a common time period. We account for this in tracker ROC-curve computations.

A related difficulty associated with the track-to-truth assignment problem often occurs as a result of track fragmentation that may be due to a target maneuver. Often, the new track will start before the old track has terminated. With both tracks classified as *true*, there is no mechanism to penalize the tracker for having overlapping-in-time tracks, which is more damaging (or confusing) to an operator than non-overlapping-in-time tracks; time-overlapping tracks can be more difficult for an operator to identify as originating from the same target.

We account for the time overlap in our performance evaluation, and include only one track segment at any time in TP_D computations. This is illustrated in Figure 6.

4.2 Intra-Ping Issue

In evaluating average track-localization error, what reference times are used? We have chosen to use ping times. There are issues associated with doing so.

Firstly, some trackers may not have track estimates available precisely at these times. For instance, if intra-ping effects are properly accounted for, we know that ensonification times differ from ping times, and accordingly so will track update times [17].

Secondly, even if all contacts are treated as ping-time observations, there may be multiple contact files with the same ping time. Thus, some trackers will have more than one track update at the same ping time. For consistency among all trackers, it is best to use only the final track update in evaluating track localization error.

4.3 Real-Time vs. Batch

Some trackers are based on batch processing of contact data. While the metrics as we have defined them can be evaluated, we expect that these trackers will achieve good output ROC-curve performance and small *TLE*, at the cost of huge latency.

We recognize that batch tracking (unless a relatively small sliding-window approach is used) is addressing a different problem than real-time tracking: the former is applicable to surveillance or situation-assessment tasks in which results are required at the end of the surveillance period; the latter is applicable to time-critical surveillance and engagement tasks. We must exercise care in comparing performance among algorithms that address different problems and scenarios.

4.4 Fixed Tracks

How do we classify fixed clutter tracks? Since they are not target-originated, they may be classified as *false*. This implies that a tracker must be able to classify such tracks and discard them. Alternatively, one might choose to regard fixed clutter points as stationary targets that are of interest to track for subsequent assessment.

4.5 Variations on Metrics

In sections 3.1-3.2, we defined *FAR* and *TFAR* per unit time, rather than per scan of data. Indeed, the higher data rate associated with multistatics may not produce a higher *FAR* or *TFAR* per scan, but might produce higher *FAR* or *TFAR* per unit time. We have chosen the latter definition, as we believe it is more operationally relevant.

The *TFAR* metric does not reflect the time duration of false tracks. As such, it does not directly reflect the average number of false tracks with which an operator must contend: short false tracks are weighed as much as lengthy false tracks. While the average number of false tracks at any time is an interesting and potentially important metric, our *TFAR* focuses on what is perhaps a metric of even greater interest: how many operational responses per unit time are required to contend with false tracks?

Our track-localization metric (*TLE*) does not reflect errors in track velocity information. A related metric could be introduced to assess these errors.

We have defined latency in a *worst-case* sense. Some algorithms, e.g. adaptive hypothesis depth trackers, may share the same worst-case latency as a fixed depth approach, but with a smaller average latency. As mentioned in section 3.3, latency can be reinterpreted in terms of localization error.

Numerous additional metrics exist in the tracking literature. Many of these are closely connected to some of our metrics; these include the following:

- *Average time-to-confirmed*: this is reflected in the TP_D and latency metrics.
- *Probability of correct contact association*: this is a lower-level metric that does not appear to be of direct operational interest; furthermore, it is only applicable to hard data association approaches.
- *Track purity*: closely related to the previous metric; this is the percentage of contacts in a *true* track that originate from the target. (In the case of multiple targets, the target to which the track is mapped applies.)
- *Track consistency*: this is a second-order metric that evaluates the consistency between track uncertainty as reported in the state covariance matrices, and actual track localization errors. A common drawback of trackers with hard data association is that state covariance matrices tend to be optimistic, as they do not reflect data association uncertainty.

5 Conclusions and Recommendations

This Special Session at FUSION 2006 represents a first attempt to evaluate a set of multistatic sonar and radar

trackers with common datasets and common metrics. Not all the trackers in this performance study are at the same level of maturity. Thus, by necessity, the results reported are only partial. In future work, the group plans to engage in ongoing tracker development and analysis efforts that may include the following:

- Generate additional datasets with reactive targets, CW-based contacts in addition to FM-based contacts, passive sonar contacts, and multistatic radar contacts.
- Further refinement and systematic evaluation of metrics across all simulated datasets and trackers.
- The introduction of additional candidate trackers, e.g. the *probability hypothesis density* (PHD) approach [18].

6 Acknowledgements

On behalf of all the members of the MSTWG, the authors wish to thank Doug Abraham (ONR) for his efforts in setting up and supporting the activities of this group.

7 References

[1] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking*, YBS Publishing, 1995.

[2] R. Streit (ed.), *Studies in Probabilistic Multi-Hypothesis Tracking and Related Topics*, Scientific and Engineering Studies Vol. SES-98-01, NUWC Division Newport, February 1998.

[3] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, 1999.

[4] L. Stone, C. Barlow, and T. Corwin, *Bayesian Multiple Target Tracking*, Artech House, 1999.

[5] W. Blair, G. Watson, T. Kirubarajan, and Y. Bar-Shalom, Benchmark for Radar Resource Allocation and Tracking with ECM, *IEEE Transactions on Aerospace and Electronic Systems* vol. 34(4), October 1998.

[6] P. de Theije and H. Groen, Multistatic Sonar Simulations with SIMONA, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[7] D. Grimmett and S. Coraluppi, Contact-Level Multistatic Sonar Data Simulator for Tracker Performance Assessment, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[8] B. La Cour, C. Collins, and J. Landry, Multi-everything Sonar Simulator (MESS), in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[9] J. Moore and W. Blair, Practical Aspects of Multisensor Tracking, in *Multitarget-Multisensor Tracking: Applications and Advances* (Y. Bar-Shalom and W. Blair eds.), Artech House, 2000.

[10] S. Coraluppi and C. Carthel, Distributed Tracking in Multistatic Sonar, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41(3), July 2005.

[11] P. Willett and S. Coraluppi, MLPDA and MLPMHT Applied to MSTWG Data, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[12] C. Hempel Probabilistic Multi-Hypothesis Tracking for Distributed Multi-Static Active Sonar, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[13] D. Krout, W. Fox, M. Hazen, and M. El-Sharkawi, Intelligent Ping Sequencing for Multistatic Sonar Systems, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[14] P. de Theije, L. Kester, and J. Bergmans, Application of the M6T tracker to simulated and experimental multistatic sonar data, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[15] O. Gerard, S. Coraluppi, C. Carthel, and D. Grimmett, Benchmark Analysis of NURC Multistatic MHT Capability, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[16] B. La Cour, Bayesian Multistatic Tracking: Results on Simulated Data from the Multistatic Tracking Working Group, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.

[17] S. Coraluppi and D. Grimmett, Intra-ping Timing Issues in Multistatic Sonar Tracking, in *Proceedings of the 7th International Conference on Information Fusion*, June 2004, Stockholm, Sweden.

[18] O. Erdinc, P. Willett, and Y. Bar-Shalom, Probability Hypothesis Density for Multitarget Multisensor Tracking, in *Proceedings of the 8th International Conference on Information Fusion*, July 2005, Philadelphia PA, USA

Intentionally blank page

Document Data Sheet

Security Classification RELEASABLE TO THE PUBLIC		Project No.
Document Serial No. NURC-PR-2006-007	Date of Issue August 2006	Total Pages 11 pp.
Author(s) Coraluppi Stefano, Grimmett Doug, de Theije Pascal		
Title Benchmark Evaluation of Multistatic Trackers		
Abstract		
Keywords		
Issuing Organization NATO Undersea Research Centre Viale San Bartolomeo 400, 19138 La Spezia, Italy [From N. America: NATO Undersea Research Centre (New York) APO AE 09613-5000]		
Tel: +39 0187 527 361 Fax: +39 0187 527 700 E-mail: library@nurc.nato.int		